

Fig. 1 Architecture Comparison. Subordinate nodes are unshaded. In the top row, arrows represent maneuver commands. In the bottom row, arrows represent measurement transmissions. Dashed arrows in the right column represent transmissions when one of the capable nodes has failed.

have a single “combiner” spacecraft flying in formation with several “collector” spacecraft. However, from the standpoint of control, this combiner spacecraft does not have to coincide with the supervisor node in the control network.

A fundamental question for partially decentralized systems is how many supervisory nodes the system should possess. One avenue for addressing this question is to recognize that a centralized architecture with a single capable node in charge of simple subordinates might be the simplest and cheapest approach, if not for the fact that a failure of the supervisor would end the mission. On the other hand, a fully decentralized confederation of peers might be more expensive than necessary to achieve the required mission reliability, especially for larger formations. The partially decentralized, or hybrid approach, would therefore seem to be the most desirable option in many cases. This paper derives a model of the static reliability of a partially decentralized system that, coupled with an assumed cost model, allows for the optimization of the cost vs. reliability trade.

The nodes chosen as supervisors in such a design trade then each serve as centralized estimation and control nodes for corresponding subsets of the remaining subordinate nodes. Each supervisor considers the measurements of its subordinates, as well as its own data, as local information. The controls each supervisor computes based on this local information will be optimal with respect to its own “sub-formation,” but

suboptimal with respect to all the nodes collectively. To compute globally optimal controls, the supervisors may exchange data with one another, e.g. via the minimal transmission algorithm of Speyer.¹ In this sense, the supervisors act as decentralized estimation and control peers with respect to each other.

For partially decentralized architectures, the next issue is the partitioning of the global state space among the various local estimators. In most mission concepts involving satellite formation flying, the absolute positions and velocities of each spacecraft are unique and dynamically uncorrelated. Any correlations which do occur arise through common biases and process noise. It is this correlation among the local state spaces that makes local estimates suboptimal with respect to a global estimate, necessitating additional data exchange to compute the globally optimal controls. Sufficiently weak correlations among the local estimates may therefore allow simplification of the local estimators and/or the data exchange framework while preserving the global optimality of maneuvers derived from the estimates. However, there are several issues besides global optimality that may dictate the need for knowledge of the global state space. Among such issues are requirements for relative state knowledge, either for control or collision avoidance, fault detection considerations, and science/payload requirements.

Many concepts for formation flying missions rely on relative GPS as a sensor. Some works have indicated

that relatively significant correlations exist in the relative states of spacecraft estimated using GPS.^{2,3} These works assumed the presence of GPS Selective Availability (SA), and studied relatively low altitude Space Shuttle missions that fly through the top of the ionosphere, so that measurements by single frequency receivers could be expected to have significant channel-dependent range biases. These works also utilized GPS receivers with only a few channels so that it was difficult to ensure that all satellites tracked by all receivers were common, which would have tended to allow for bias cancellation in the relative states computed from their solutions. More recently, with the advent of "all-in-view" GPS receivers for space applications and the end of SA, a study⁴ has indicated that, at altitudes above the ionosphere where many currently proposed formation flying missions will orbit, channel-dependent range biases may be insignificant, so that significant correlations among the local state spaces do not arise. This paper provides some results based on covariance and monte-carlo analysis of a simplified system to show that this lack of correlation may allow simplification of the local estimators, while preserving the global optimality of the maneuvers commanded by the supervisors.

The remainder of this paper first describes the reliability/cost trade, including assumptions that simplify the analysis, analytic models for the reliability, a simple cost model, and results of some case studies. The next section reviews the minimal data transmission requirements for globally optimal control for the partially decentralized system, and examines how the navigation architecture may be simplified if the correlations among the local state spaces are not significant. The final section provides a summary and conclusions, with suggestions for future work.

Reliability vs. Cost Trade

One of the inferred benefits of decentralized control approaches is that they will be more reliable due to their non-hierarchical and inherently parallel structure. The purpose of the study reported here was to perform some simple static reliability calculations to investigate this claim, and in particular to find the most cost-effective architecture that will satisfy a specified mission reliability requirement.

An additional complexity is that any node may be built with redundant "strings" so that it can tolerate one or more failures. The analysis below incorporates this concept, and it appends a number to the subscripts of q , the probability of failure, and c , the system cost, to indicate the number of redundant strings per node. For example, q_{c2} is the failure probability of two-string capable node, and c_{c2} is its associated cost. Nominally, parallel redundancy would imply that $q_{cn} = q_{c1}^n$. However, adding redundancy also increases weight and complexity, and according

to Reference 5, experience has shown that there are diminishing returns to increasing reliability through redundancy. Reference 5 addresses this through parametric power laws whose exponents depend on either weight or cost ratios. To avoid the need for baseline costs or weights which these models require, this study penalizes the use of additional strings so that

$$q_{cn} > q_{c1}^n. \quad (1)$$

In particular, this study assumes that

$$q_{in} = q_{i1}^{1+\log n}, \quad (2)$$

where $i = c, s$. Figure 2 compares the exponents of Eqs. (1) and (2). Although Eq. (2) is not em-

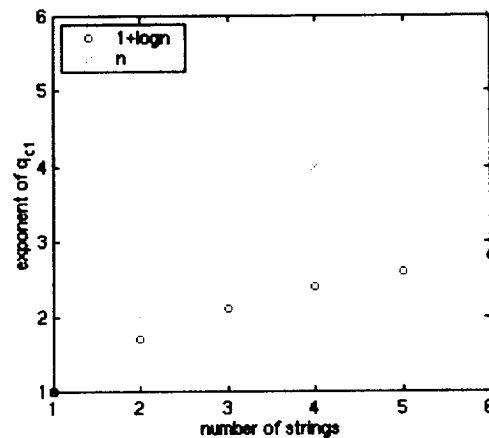


Fig. 2 A model of diminishing returns to redundancy: the power to which the single-string failure probability, q_{c1} , is raised as additional strings are added, in comparison with the ideal case in which the exponent is n .

pirical, Figure 2 illustrates that the failure probability decreases less significantly after the third string, which is consistent with the observation that very few aerospace systems have more than three strings.

This analysis does not consider "design flaws," which would be common failure modes to all similar nodes. This work assumes that each failure is an independent event. This analysis also does not consider that there are any improvements in reliability due to a learning curve as multiple copies of a design are produced. Each node of a similar type costs the same to produce or replace as any other, whether or not several or many of the same type have already been built.

In estimating costs, this work considers only the initial cost to build the formation. One could consider as well expected replacement costs; however, unless the failure probabilities are quite large, the expected value of this cost may be quite small in comparison to the deterministic initial build cost. One must also decide how many "rounds" of replacements to consider, as well as how to handle mean time between failures, etc.

Static Reliability Model

Let r_M represent the mission reliability. As stated above, the purpose of this study is to find the most cost-effective architecture that will satisfy a specified mission reliability, e.g. $r_M = 0.95$. The reliability is the probability that the system works, or in the present case, that at least some minimum number, k_{min} , out of the k spacecraft control systems operate correctly:

$$r_M = \Pr(\geq k_{min} \text{ out of } k \text{ work}) \quad (3)$$

Centralized Architecture

Figure 3 is a reliability diagram for the centralized architecture. In this diagram, each switch represents a node or spacecraft, and an open switch indicates a total failure of that node's control system. For the mission to operate, k_{min} of the switches must be closed. The probability of each switch being closed is noted above it in the figure.

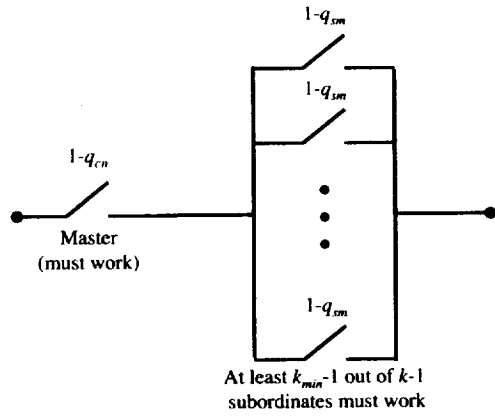


Fig. 3 Centralized Architecture Reliability Diagram

For the centralized architecture, regardless of what happens with the rest of the nodes, if the master fails, a mission failure occurs. Given that the master works, at least $k_{min} - 1$ subordinates must work in order that k_{min} total nodes function correctly. Therefore, the mission reliability is the probability that both the master works and that at least $k_{min} - 1$ out of $k - 1$ subordinates work:

$$r_{Mc} = \Pr(\text{master works}) \times \Pr(\geq k_{min} - 1 \text{ of } k - 1 \text{ subs. work}) \quad (4)$$

As Figure 3 indicates, the probability that the n -string master node works is $1 - q_{cn}$. For the subordinate nodes, the binomial distribution gives the probability that *exactly* j out of $k - 1$ subordinates operate:

$$P_{k-1}^{sm}(j) = \binom{k-1}{j} (1 - q_{sm})^j q_{sm}^{k-1-j} \quad (5)$$

The probability that *at least* $k_{min} - 1$ out of $k - 1$ subordinates operate is then

$$\Pr(\geq k_{min} - 1) = \sum_{j=k_{min}-1}^{k-1} P_{k-1}^{sm}(j) \quad (6)$$

From Eqs. (4) and (6), the total reliability for the centralized scenario is therefore

$$r_{Mc} = (1 - q_{cn}) \sum_{j=k_{min}-1}^{k-1} P_{k-1}^{sm}(j). \quad (7)$$

Decentralized Architecture

Figure 4 is a reliability diagram for the decentralized architecture. As the figure indicates, this case is structurally identical to the subordinate node network in the centralized architecture, but it does not rely on the operation of a master node. Therefore, the decentralized system's reliability is

$$r_{Md} = \sum_{i=k_{min}}^k P_k^{cn}(i), \quad (8)$$

where

$$P_k^{cn}(i) = \binom{k}{i} (1 - q_{cn})^i q_{cn}^{k-i} \quad (9)$$

is the probability that exactly i out of the k peers operate correctly.

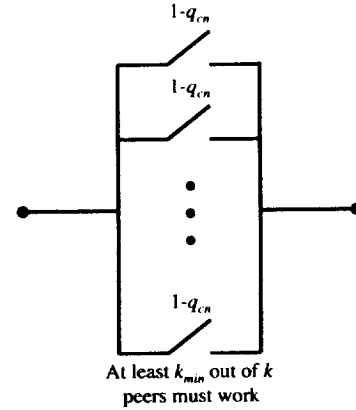


Fig. 4 Decentralized Architecture Reliability Diagram

Partially Decentralized Architecture

Figure 5 is a reliability diagram for the partially decentralized or hybrid architecture. For the mission to operate, k_{min} of the nodes must operate, but of those k_{min} operating nodes, at least one out of the ℓ capable nodes must work in order to serve as the master. Therefore, the reliability of the hybrid architecture is

$$r_{Mh} = \Pr(\geq 1 \text{ of } \ell \text{ cap. nodes work}) \times \Pr(\text{"enough" subs. work}) \quad (10)$$

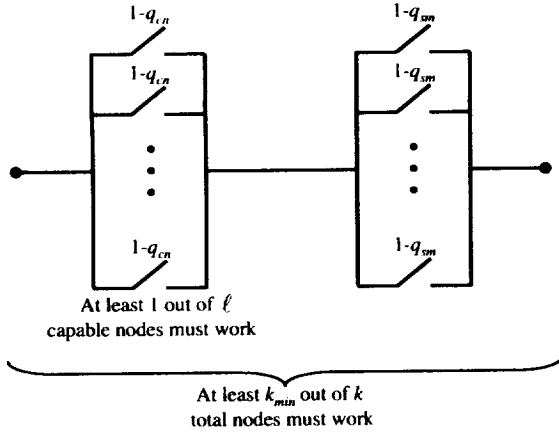


Fig. 5 Decentralized Architecture Reliability Diagram

There are several ways of looking at the problem of computing the probability Eq. (10) describes. The approach this paper takes is to consider the following four distinct cases:

1. $k - \ell \geq k_{min} - 1$ and $\ell < k_{min}$. In this case there are not enough capable nodes to make up the required complement of k_{min} total nodes, so at least some subordinate nodes must function.
2. $k - \ell \geq k_{min} - 1$ and $\ell = k_{min}$. In this case there are exactly enough capable nodes so that if they all work, no subordinates need to operate.
3. $k - \ell \geq k_{min} - 1$ and $\ell > k_{min}$. In this case there are more than enough capable nodes so that if they all work, no subordinates need to operate.
4. $k - \ell < k_{min} - 1$ and $\ell > k_{min}$. In this case there are too few subordinate nodes, so that more than one capable node must work in order to achieve k_{min} total operating nodes.

Eqs. (11)–(13) capture all four cases:

$$r_{Mh} = \sum_{i=i_{min}}^{i_{max}} P_{\ell}^{cn}(i) \sum_{j=k_{min}-i}^{k-\ell} P_{k-\ell}^{sm}(j) + \sum_{i=i_{max}+1}^{\ell} P_{\ell}^{cn}(i) \quad (11)$$

where

$$i_{min} = \begin{cases} k_{min} - (k - \ell) & \text{if } k - \ell < k_{min} - 1 \\ 1 & \text{otherwise} \end{cases} \quad (12)$$

$$i_{max} = \begin{cases} k_{min} - 1 & \text{if } \ell > k_{min} \\ \ell & \text{otherwise} \end{cases} \quad (13)$$

In Eq. (11), all summations only run forward; if the upper index is greater than the lower index, the summation is zero by assumption.

When $\ell \leq k_{min}$, Cases 1 and 2 apply. If there are too few capable nodes to make up the required complement of k_{min} total nodes, Eq. (13) stops the first summation over i in Eq. (11) at ℓ , and the nested summation over j always starts with enough subordinate nodes to make up the deficit of capable nodes.

When $\ell > k_{min}$, Cases 3 and 4 apply. In both of these cases, the second summation over i in Eq. (11) captures the probability that more than k_{min} capable nodes may operate. Eq. (13) generates these cases by stopping the first summation and starting the second summation when there are no more subordinate nodes available to include under the nested summation over j . When also $k - \ell < k_{min} - 1$, Case 4, where there are too few subordinate nodes, applies. Eq. (12) generates this case by starting the first summation over i in Eq. (11) with enough capable nodes to achieve k_{min} total operating nodes.

Cost Model

The cost to build the centralized architecture is the cost of a k -node formation, using an n -string master and $k - 1$ m -string subordinates:

$$c_{co} = c_{cn} + (k - 1)c_{sm}. \quad (14)$$

The cost to build the fully decentralized architecture is the cost of a k -node formation using n -string peers:

$$c_{do} = k \cdot c_{cn}. \quad (15)$$

The cost to build a partially decentralized architecture is the cost of ℓ capable nodes with n strings each, plus the cost of $k - \ell$ subordinate nodes with m strings each:

$$c_{co} = \ell c_{cn} + (k - \ell)c_{sm}. \quad (16)$$

This study presumes an inverse correlation between reliability and cost:

$$c_{c1} \propto q_{c1}^{-1} \quad (17)$$

$$c_{s1} \propto q_{s1}^{-1} \quad (18)$$

This model is consistent with the concavity of the models of Reference 5. To scale the various costs, assume that a single-string, capable node having a failure probability of 0.1 costs one unit. Assume that a comparably reliable subordinate node costs half as much, $c_{s1} = 0.5c_{c1}$, and that adding a comparable extra string to a node costs three-quarters as much as the node itself, $c_{is} = 0.75c_{i1}$. As a consequence of these further assumptions,

$$c_{c1} = 0.1q_{c1}^{-1} \quad (19)$$

$$c_{s1} = 0.05q_{s1}^{-1}, \quad (20)$$

Case Studies

This section examines four sets of mission parameters, k and k_{min} , corresponding to “small” and “large”

formations. All cases have a specified requirement of an overall reliability, r_M . To achieve this requirement, one may vary the design parameters q_{c1} , q_{s1} , n and m , and select a centralized, decentralized, or a partially decentralized hybrid architecture, with the goal of achieving the minimum mission cost.

For the decentralized architecture, one may choose a design point for each number of strings, n , that exactly satisfies the reliability requirement, which may be found by solving Eq. (8) for q_{cn} ; choosing the largest appropriate root will minimize the cost. Eq. (2) gives the resulting q_{c1} , and one may select the number of strings that minimizes the build cost, Eq. (15).

For the centralized architecture, one may vary both q_{c1} and q_{s1} to achieve the reliability requirement. Therefore to find the minimum cost, one must minimize Eq. (14) with respect to either q_{c1} or q_{s1} , subject to Eqs. (7) and (2). One may then select the number of strings that minimize the costs among the various possible combinations of n and m . This study limited the number of strings to a maximum of six. As a result of the diminishing returns to redundancy Eq. (2) imposes, the optimal number of strings never exceeded five, and the differences in cost between three, four, and five string optima were only a few percent. When preliminary studies did not impose diminishing returns to redundancy, unrealistically large numbers of strings often arose as the optima.

One may optimize the hybrid architecture in a similar fashion to the centralized architecture, with the additional degree of freedom to choose the number of capable nodes.

Tables 1 through 5 summarize the results of following the procedure above. The mission parameters all reflect actual or proposed missions. The notation (k, k_{min}, ℓ) is a handy shorthand for referring to each case.

The *StarLight* (ST-3) mission⁶ inspires the $(2, 2, \ell)$ and $(3, 2, \ell)$ cases. This mission originally had a $(3, 2, 1)$ configuration, but was trimmed to a $(2, 2, 1)$ configuration to save costs. The Magnetospheric Multi-Scale,⁷ or MMS, and Stellar Interferometer,⁸ or SI, missions inspire the $(6, 4, \ell)$ case and the $(30, 20, \ell)$ case, respectively. Note that for the $(2, 2, \ell)$, the hybrid architecture does not apply, since $\ell = 1$ corresponds to the centralized case and $\ell = 2$ corresponds to the fully decentralized case.

Consider first the scenario without multiple strings, which Tables 1 and 2 present. For a 95% reliability requirement, Table 1 shows that the minimum cost architecture, under the assumptions of this study, is always the hybrid or partially decentralized option. It is interesting that only $\ell = 4$ capable nodes give the minimum cost for the large $(30, 20, \ell)$ formation. The fully decentralized architecture is more cost-effective than the centralized for small and medium formations that can tolerate a few failures, i.e. the $(3, 2, \ell)$ and $(6, 4, \ell)$

cases. The centralized architecture is cheaper than the fully decentralized for the minimal $(2, 2, \ell)$ formation and for the large $(30, 20, \ell)$ formation. For a 99% reli-

Table 1 Results for mission reliability $r_M = 0.95$, single-strings only

	k	2	3	6	30
	k_{min}	2	2	4	20
Dec.	q_{c1}	.025	.135	.153	.221
	c_M	7.90	2.21	3.91	13.6
	ℓ_{opt}	-	2	3	4
Hyb.	q_{c1}	-	.154	.185	.266
	q_{s1}	-	.101	.123	.212
	c_M	-	1.79	2.84	7.65
Cen.	q_{c1}	.030	.041	.038	.036
	q_{s1}	.021	.095	.113	.191
	c_M	5.76	3.47	4.82	10.3

ability requirement, Table 2 shows that the minimum cost architecture remains the partially decentralized, with the same numbers of capable nodes for each case as for the 95% reliability requirement. The fully decentralized architecture is more cost-effective than the centralized for all but the $(2, 2, \ell)$ cases.

Table 2 Results for mission reliability $r_M = 0.99$, single-strings only

	k	2	3	6	30
	k_{min}	2	2	4	20
Dec.	q_{c1}	.005	.059	.009	.176
	c_M	39.9	5.09	7.08	17.1
	ℓ_{opt}	-	2	3	4
Hyb.	q_{c1}	-	.067	.103	.196
	q_{s1}	-	.043	.068	.169
	c_M	-	4.12	5.13	9.74
Cen.	q_{c1}	.006	.009	.009	.009
	q_{s1}	.004	.034	.052	.137
	c_M	29.1	14.2	16.3	21.7

For the single-string cases, it appears that the fully decentralized architecture's lower node-level reliability requirements outweigh its additional node-level costs in most cases. The exceptions are the minimal $(2, 2, \ell)$ configurations, and the lower (95%) reliability $(30, 20, \ell)$ case. The hybrid partially decentralized architecture allows the designer the flexibility to utilize a cheaper intermediate architecture than either.

For the $(2, 2, \ell)$ case, the fully decentralized peer nodes' reliability requirement is of the same order as the master and subordinate nodes' reliability requirements (about midway between). Since its cost is therefore of the same order as the cost for two master nodes, it is more expensive than the centralized architecture.

For the $(30, 20, \ell)$ case, the fully decentralized peer nodes' reliability requirement is lower than the centralized nodes'. However, many relatively expensive peer nodes are required for this configuration. When

the overall mission reliability requirement is higher, it is cheaper to achieve using less reliable but relatively more expensive capable nodes in a decentralized network. As the overall mission reliability drops, the converse appears to be true. The fact that only four capable nodes allows the hybrid architecture to achieve the same reliability at much lower cost indicates some sense of the inefficiency of the fully decentralized architecture for this large formation.

Next, consider the effects of allowing for multiple strings. Tables 3 through 5 present these results. At

Table 3 Results for mission reliability $r_M = 0.95$

	k	2	3	6	30
	k_{min}	2	2	4	20
	n	3	2	2	2
Dec.	q_{c1}	.173	.307	.330	.410
	c_M	2.88	1.71	3.18	12.8
	ℓ_{opt}	-	2	2	4
	n	-	2	2	1
	m	-	2	2	2
Hyb.	q_{c1}	-	.338	.318	.316
	q_{s1}	-	.244	.308	.388
	c_M	-	1.39	2.24	7.14
	n	3	3	3	3
	m	3	2	2	2
Cen.	q_{c1}	.191	.210	.198	.182
	q_{s1}	.153	.276	.297	.391
	c_M	2.12	1.83	2.74	7.86

the 95% mission reliability level, Table 3 shows that the hybrid architecture is again the most cost effective. The fully decentralized architecture is only favored over the centralized architecture in the $(3, 2, \ell)$ case, and this advantage is less than 10%. It appears that at this mission reliability level it is more cost-effective to add reliability via extra strings on less expensive subordinate nodes than via a fully decentralized architecture.

At the 99% and 99.7% mission reliability levels, Tables 4 and 5 show results that are more in line with the single-string results. Again, the hybrid architecture is clearly the most cost effective. For the higher mission reliability requirements, the minimal $(2, 2, \ell)$ and large $(30, 20, \ell)$ formations are more cost-effectively accomplished with a centralized than a fully decentralized design, but the fully decentralized approach appears to be cheaper than the centralized for the intermediate configurations. The rationale for these results appears to be similar to the single-string cases. That is, for the $(2, 2, \ell)$ case, there is no reliability advantage for the fully decentralized case to combat its extra node-level cost, and for the $(30, 20, \ell)$ case, the large number of more expensive peer nodes required outweighs their individually lower reliability requirement.

The bottom of Table 4 contains an additional tier of rows that present the results of a direct 5000-case monte carlo simulation of the systems that Figures 3,

Table 5 Results for mission reliability $r_M = 0.997$

	k	2	3	6	30
	k_{min}	2	2	4	20
	n	4	3	2	2
Dec.	q_{c1}	.066	.194	.181	.327
	c_M	9.91	3.87	5.80	16.1
	ℓ_{opt}	-	2	3	4
	n	-	3	2	2
	m	-	3	3	2
Hyb.	q_{c1}	-	.212	.220	.334
	q_{s1}	-	.156	.206	.320
	c_M	-	3.16	4.20	9.21
	n	5	5	5	5
	m	5	3	3	2
Cen.	q_{c1}	.090	.100	.098	.096
	q_{s1}	.074	.168	.220	.304
	c_M	7.13	5.50	6.93	12.5

4, and 5 depict, using the values for q_{c1} , q_{s1} , ℓ , n , and m from Table 4. If the analysis used in solving for these parameters were incorrect, the monte carlo simulation should produce estimates of the mission reliability, r_M , that differ at a statistically significant level from the 99% mission reliability value used in deriving them. As the two bottom rows of Table 4 indicate, this is not the case, since the 95% confidence limits on r_M bound the assumed value of 99% mission reliability. This implies that there is only a 5% probability that the true value of r_M lies outside the indicated ranges.

Partitioning of the Global State Space, Data Transmission, and Implications for GPS Relative Navigation

Regardless of the number of capable nodes a particular design utilizes, unless it is a fully centralized architecture, an important next step in the design process is determining how to partition the global state space among the various local estimators. Willsky, et al.⁹ showed that the state space for the local estimators need only satisfy the necessary and sufficient condition that there exist matrices M^j , where $j = 1, 2, \dots, k$, such that

$$C^j = H^j M^j, \quad (21)$$

where C^j is the map from the global state space to the local measurement \mathbf{y}^j , and H^j is the map from the local state space to the local measurement. Among the many possibilities Eq. (21) allows is that the global state space could consist of the earth-centered inertial states of all the spacecraft, while the local state spaces could consist of the earth-centered ("absolute") states of each spacecraft, expressed in any earth-centered coordinate system, along with the relative states of all the other spacecraft with respect to the local spacecraft, expressed in a local spacecraft-centered coordinate system. An even simpler approach would have

Table 4 Results for mission reliability $r_M = 0.99$

	k	2	3	6	30
	k_{min}	2	2	4	20
	n	4	2	2	2
Dec.	q_{c1}	.109	.188	.233	.358
	c_M	5.98	2.80	4.51	14.7
	ℓ_{opt}	-	2	3	4
	n	-	2	2	2
	m	-	3	2	2
Hyb.	q_{c1}	-	.214	.270	.386
	q_{s1}	-	.200	.195	.349
	c_M	-	2.26	3.29	8.34
	n	4	4	4	4
	m	4	2	2	2
Cen.	q_{c1}	.119	.130	.126	.122
	q_{s1}	.097	.166	.205	.338
	c_M	4.41	3.55	4.71	10.2
5000-case monte carlo check; 95% confidence limits on r_M :					
Dec.		[.984,.990]	[.987,.992]	[.987,.992]	[.988,.993]
Hyb.		-	[.986,.992]	[.986,.991]	[.989,.994]
Cen.		[.986,.992]	[.989,.994]	[.988,.993]	[.984,.990]

all the spacecraft, including the subordinates, locally estimate only their own absolute states. The capable spacecraft could then difference transmissions of their subordinates' absolute states with their own absolute states to arrive at the relative states which the formation control law may require.

The potential problem with the simple "state vector differencing" approach, is that it fails to capture correlations among the local state spaces that may arise through correlated states (e.g. biases common to multiple nodes), correlated process noise, and/or correlated sources of initialization data. A centralized estimator receiving all the measurement data and estimating the full global state space would capture these correlations correctly. Speyer¹ proposes a more efficient approach for decentralized systems, in which local estimators at each node process only local data, generating locally optimal but globally suboptimal state estimates. A brief summary of this algorithm highlighting topics relevant to the present work follows, along with a discussion of potential simplifications of the local state spaces, and a simple illustrative example.

Minimum Data Transmission

In Speyer's algorithm, each node partitions its state space into a "control dependent" and "data dependent" partition. The control dependent partitions receive only the initial condition and control update information, while the data dependent partition is updated with a local Kalman filter. The local nodes also maintain an additional "data vector" \mathbf{h}_i^j , at each epoch i , with the following recursion:

$$\mathbf{h}_i^j = F_i \mathbf{h}_{i-1}^j + G_i^j \bar{\mathbf{x}}_i^{Dj}, \quad \mathbf{h}_1^j = 0 \quad (22)$$

where $\bar{\mathbf{x}}_i^{Dj}$ is the *a priori* estimate of the data dependent partition of the state, and F_i and G_i^j are computed using the local and global system parameters and covariances. Note that this recursion also requires that each node maintain a local copy of the global covariance matrix. Now, when it is time to compute a globally optimal state estimate, the capable nodes compute and exchange the following vectors

$$\beta_i^j = P_i (\hat{P}_i^j)^{-1} \mathbf{h}_i^j \quad (23)$$

where P_i represents the *a posteriori* global covariance matrix and \hat{P}_i^j represents the *a posteriori* local covariance matrix. All the received β_i^j are then summed with each capable node's local copy of the control dependent partition of the global state space to arrive at the globally optimal state estimate. Although the spacecraft may employ any control, if they use a linear-quadratic type of control, they may reduce data transmission requirements further. In this case, they reconstruct the globally optimal control without reconstructing the globally optimal state, by exchanging pairs of vectors $\alpha_i^{j\ell}$ between every two nodes j and ℓ , where $\alpha_i^{j\ell}$ has only the dimension of the control.¹

In order to propagate the state space of the other nodes, each node needs to know the other nodes' controls as well the state information. In theory, so long as every capable node has the schedule of maneuvers and knowledge of the control parameters at every other node, the nodes need not exchange the controls, since each node can compute the controls that the others perform. In practice, the actual controls may differ from the computed controls, so the nodes should exchange their maneuvers as well. This control information may also be useful in fault detection.

References 10 and 11 studied the optimal fully decentralized formation control architecture using Speyer's data transmission framework. In these works, all of the states present in the full global state vector were maintained by each local estimator, regardless of their local observability. These works did not explicitly study the importance of the correlations among the local position and velocity states. However, Reference 10 studied a simplified case with two-body dynamics and absolute position measurements similar to GPS point solutions, with no measurement biases, so few if any correlations should have existed. Reference 11 studied a somewhat more complex case, adding J_2 to the gravity model as well as measurements of the relative position vectors between the satellites, which could have introduced some intersatellite state correlation.

The Role of Correlation in Simplifying the Local Estimators

If Reference 10 is representative of realistic applications in terms of the role of correlations, as Reference 4 indicates it may be for some relative GPS applications, then the data transmission framework is not required, potentially offering significant simplification of the local estimators. In this case, all nodes could dispense with the burden imposed by Eqs. (22) and (23) and they could estimate and exchange only their own absolute positions and velocities. Most GPS receivers for space applications already perform this function, although as Reference 12 describes, not all have done as good a job of this as one might expect. Otherwise, it may be important to implement the data exchange framework.

In the context of the partially decentralized architecture, the minimal data transmission approach forms an efficient framework for the peer-to-peer decentralized control in which the capable nodes engage. If correlations are significant, the capable nodes would exchange the β_i^j 's at scheduled maneuver epochs in order to optimally estimate the entire global state space for use in computing the maneuvers, as Figure 6(a) shows. Nominally, the subformations would operate as centralized frameworks with the subordinates sending measurements to their respective capable node for processing. The minimal data transmission framework could also serve as a decentralized estimation scheme at the level of the sub-formations, as Figure 6(b) shows. Here each node would be responsible for its own local state estimation, and the capable nodes would perform a fusion of the local estimates, using Eq. (23). At the sub-formation level, the minimal data transmission approach could substantially reduce data transmission requirements over the centralized approach of sending all the measurement data. If correlations are significant, sending measurements as in Figure 6(a) instead of employing the minimal data transmission framework as in Figure 6(b) will generally impose a higher

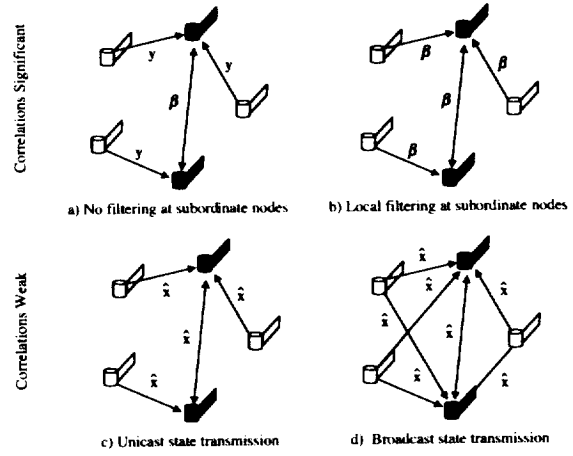


Fig. 6 Data Transmission Options for the Partially Decentralized Architecture

data transmission burden as a cost of simpler local processing. If correlations are weak, and the network exchanges only state estimates, then the unicast network of Figure 6(c) requires a simpler communications topology than the broadcast network of Figure 6(d), but may induce additional delays since the capable nodes must act as relays.

Example

To examine these issues, consider the following simple example problem. The example system consists of two nodes, denoted by $j = 1, 2$, each of which makes bias-corrupted measurements of a unique position:

$$y_i^j = r_i^j + b_i^j + v_i^j, \quad (24)$$

where $E[v_i^j] = 0$, $E[v_i^j(v_h^j)^T] = R_i^j \delta_{ih}$. The position r_i^j is a random walk:

$$\dot{r}^j(t) = u^j(t), \quad (25)$$

where $E[u^j(t)] = 0$, $E[u^j(t)(u^j(\tau))^T] = Q(t)^j \delta(t, \tau)$. The bias and initial position are random constants:

$$b_1^j = 0, \quad (26)$$

where $E[b_1^j] = 0$, $E[b_1^j(b_1^j)^T] = P_1^{bj}$, and $E[r_1^j] = 0$, $E[r_1^j(r_1^j)^T] = P_1^{rj}$. The example compares a centralized estimator processing all the measurements to a suboptimal decentralized estimator that processes only local measurements, and performs no information exchange. The goal of the example is to estimate the relative state, $r_1^2 - r_1^1$, by differencing the absolute state estimates, and its covariance matrix. Both filters estimate the position and bias as states. The centralized filter maintains the full 4×4 covariance, while the local suboptimal filters only maintain the 2×2 covariances corresponding to their local estimates. Therefore, only the centralized filter has access to the cross-covariance between position 1 and position 2, which appears in

the calculation of the relative state covariance,

$$\hat{P}_i^{rel} = \hat{P}_i^{r1} + \hat{P}_i^{r2} - \hat{P}_i^{12} - (\hat{P}_i^{21})^T. \quad (27)$$

The cross-covariance \hat{P}_i^{12} contains the correlations that generate the need for data exchange. If these correlations are not significant, then $\hat{P}_i^{12} \approx 0$, and the suboptimal local filters should be nearly optimal at estimating the relative states, even without data exchange.

The example consists of three cases. In the first, there are no correlations among any of the random variables, i.e. $E[r_i^1(r_i^2)^T] = 0$, $E[b_i^1(b_i^2)^T] = 0$, $E[v_i^1(v_i^2)^T] = 0 \forall i$, $E[w_i^1(w_i^2)^T] = 0 \forall i$. In the second case, the biases are perfectly correlated i.e. $b_i^1 = b_i^2$. In the third case, the process noises are perfectly correlated, i.e. $w_i^1 = w_i^2 \forall i$. The first case abstracts a scenario similar to that studied in Reference 4, in which the spacecraft are at altitudes above the ionosphere, and there is no SA, so that channel-dependent range biases may be insignificant, and significant correlations among the local state spaces do not arise. Figure 7(a) illustrates the results of covariance and monte carlo analysis of this case. The second case is an extreme example of the kind of situation that occurs when two receivers track the same GPS satellite, and the measurements they make contain the same SA and/or ionosphere bias, as occurred in some of the results in References 2 and 3. Figure 7(b) illustrates the results of covariance and monte carlo analysis of this case. The third case is an extreme example of the kind of situation that occurs when the filters use process noise to account for missing dynamics terms, such as higher-order gravity. If the spacecraft are close to each other, the same acceleration “noise” due to the neglected dynamics will affect both of them. Figure 7(c) illustrates the results of covariance and monte carlo analysis of this case.

In the first case where there is no correlation, the local suboptimal filters’ covariance closely matches the actual covariance of the 30 monte carlo cases, and the performance appears to be indistinguishable from that of the centralized filter. In this case, then, data exchange is superfluous. The other two cases exhibit significant differences in performance between the centralized and local filters. In the case of the common bias, the local filters perform somewhat worse than the centralized, and substantially overestimate their covariances. In the equal process noise case, the centralized filter’s performance is much better than that of the local filters, but the local filters at least correctly estimate their covariances. If either of these cases is representative of the actual scenario, then data exchange will likely be a requirement.

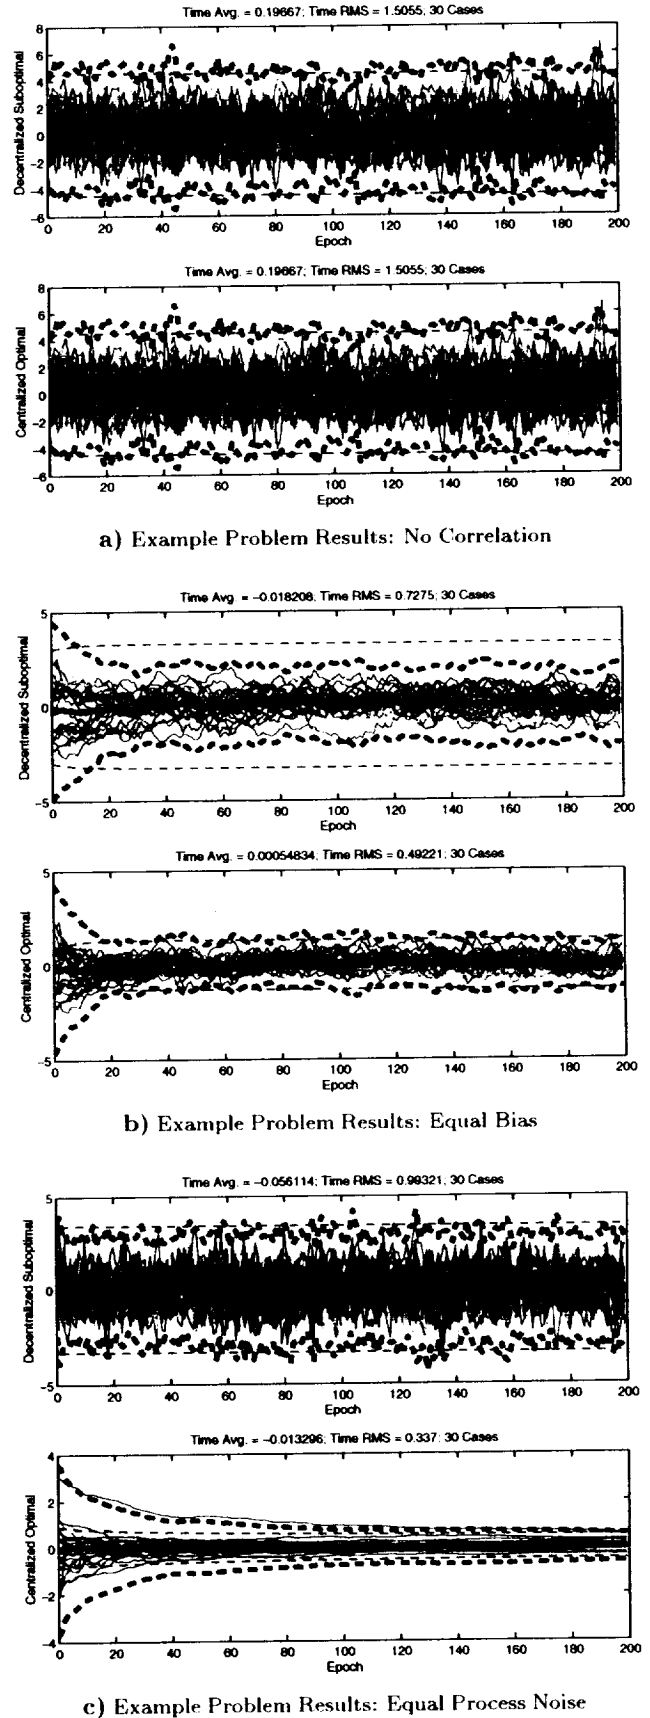


Fig. 7 Thick solid lines are ensemble means. Light dashed lines are filters’ 3 σ bounds. Heavy dashed lines are ensemble 3*RMS bounds.

Summary, Conclusions, and Future Directions

This study has presented analytic models for the reliabilities of centralized, fully decentralized, and hybrid, partially decentralized control architectures. Coupled with a cost model, these models may be used to find the most cost-effective architecture, from the standpoint of the guidance, navigation, and control function, for a given mission. In case studies of four architectures proposed for satellite formation flying missions, along with an assumed cost model, the partially decentralized architecture always achieved the required mission-level reliability at much lower cost than either the fully decentralized or the centralized approaches. The number of capable nodes that achieved the lowest cost for the hybrid architecture tended to be higher when the fully decentralized architecture was more favorable than the centralized, and vice versa.

The study found that the fully decentralized architecture is more cost-effective than the centralized for smaller formations that can tolerate a few failures, especially for higher mission-level reliability requirements. One can attribute this to the graceful degradation capability inherent in a decentralized architecture better taking advantage of the surplus nodes in failure scenarios.

The centralized architecture is favorable over the fully decentralized for a minimal formation of only two satellites. In this case, the fully decentralized architecture requires similar node-level reliability to the centralized case, so its requirement for two capable nodes makes it less cost-effective.

For larger formations, the additional cost of the many capable nodes the fully decentralized architecture requires again may put it at a disadvantage in comparison to the centralized approach, especially when the mission-level reliability requirement is lower. For such cases, the hybrid architecture is especially appealing, since with only a few capable nodes, such a configuration can overcome the single-point failure disadvantage of the centralized architecture, while avoiding the additional cost of making every node a fully capable peer.

Nominally, the nodes chosen as supervisors each serve as centralized estimation and control nodes for their respective subsets of subordinate nodes, and the supervisors act as decentralized estimation and control peers with respect to each other. However, for many currently proposed formation flying missions that will use GPS, if significant correlations among the local state spaces do not arise, simplifications of the local estimators that preserve the global optimality of the maneuvers commanded by the supervisors may be possible.

Further work should bring in the concept of time, so that failure rates and mean time between failures can

be compared with desired mission lifetimes. Future work should also consider cost models based on actual design and build experience. More detailed relative navigation studies also need to be performed in the context of the centralized/decentralized architecture trade.

Acknowledgments

Landis Markley suggested the idea of studying partially decentralized architectures. Teddie Brinkley provided independent checks of some of the results.

References

- ¹Speyer, J. L., "Computation and Transmission Requirements for a Decentralized Linear-Quadratic-Gaussian Control Problem," *IEEE Transactions on Automatic Control*, Vol. AC-24, No. 2, April 1979, pp. 266-269.
- ²Montez, M. N. and Zyla, L. V., "Use of Two GPS Receivers in Order to Perform Space Vehicle Orbital Rendezvous," *Proceedings of the 6th International Technical Meeting of the Satellite Division (ION-GPS)*, The Institute of Navigation, Alexandria, VA, Sept. 1993, pp. 301-312.
- ³Schiesser, E., Brazzel, J. P., Carpenter, J. R., and Hinkel, H. D., "Results of STS-80 Relative GPS Navigation Flight Experiment," *SPACEFLIGHT MECHANICS 1998*, Vol. 99 of *Advances in the Astronautical Sciences*, Univelt, San Diego, CA, Feb. 1998, pp. 1317-1334.
- ⁴Gramling, C., Carpenter, J. R., Long, A., Kelbel, D., and Lee, T., "Autonomous Relative Navigation for Formation-Flying Satellites Using GPS," *15th International Symposium on Space Flight Dynamics*, CNES, Toulouse, France, June 2000, Paper MS00/18.
- ⁵Hecht, H., "Reliability During Space Mission Concept Exploration," *Space Mission Analysis and Design*, edited by J. R. Wertz and W. J. Larson, Kluwer Academic Publishers, Boston, 1991, pp. 640-656.
- ⁶"StarLight," <http://starlight.jpl.nasa.gov/>, Oct. 2001.
- ⁷"Magnetospheric Multiscale," <http://stp.gsfc.nasa.gov/missions/mms/mms.htm>, Oct. 2001.
- ⁸"The GSFC Stellar Interferometer (SI) Homepage," <http://hires.gsfc.nasa.gov/si/>, July 2001.
- ⁹Willsky, A. S., Bello, M. G., Castanon, D. A., Levy, B. C., and Verghese, G. C., "Combining and Updating of Local Estimates and Regional Maps Along Sets of One-Dimensional Tracks," *IEEE Transactions on Automatic Control*, Vol. AC-27, No. 4, Aug. 1982, pp. 799-813.
- ¹⁰Carpenter, J. R., "A Preliminary Investigation of Decentralized Control for Satellite Formations," *Aerospace Conference Proceedings, CD-ROM Edition*, IEEE Publications, Piscataway, NJ, March 2000, Paper 264.
- ¹¹Carpenter, J. R., "Decentralized Control of Satellite Formations," *International Journal of Robust and Nonlinear Control*, 2002, (to appear).
- ¹²Carpenter, J. R. and Schiesser, E. R., "Semimajor Axis Knowledge and GPS Orbit Determination," *NAVIGATION: Journal of The Institute of Navigation*, Vol. 48, No. 1, Spring 2001, pp. 57-68, also AAS Paper 99-190, Feb., 1999.